A Multi-Agent Environment for Department of Defense Distribution

Laurence Glicoes
LMI
2000 Corporate Road
McLean, Virginia 22101

Richard Staats
LMI
2000 Corporate Road
McLean, Virginia 22101

Michael N. Huhns
MCC
3500 West Balcones Center Drive
Austin, Texas 78759

Abstract

The United States Department of Defense (DoD) requires an effective, economic method for utilizing available distribution system to move its personnel, equipment and supplies in support of military operations world wide. Recent reductions in the DoD budget have placed a premium on leveraging technologically innovative solutions to accomplish this requirement. This paper examines the integration of cooperative autonomous computational agent technology with low cost satellite communications capability. Under this concept, Intelligent Agents (IA) would be developed and integrated into the spectrum of transportation actions DoD wide. The IA would be divided into two categories, static (attached to intermodal sites) and mobile (attached to shipments). The IA act as economic competitors in routing the shipments through the DoD transportation network. The result being effective and efficient transportation of goods and personnel for both routine operations and unforeseen contingencies. The global communication system offered by the satellites would be used to track shipment status and continually update the shared intermodal knowledge base.

1. Introduction

After the Gulf War, the Department of Defense (DoD) embarked upon a number of logistics initiatives directed toward eliminating problems in the movement of personnel and equipment. Total Asset Visibility (TAV), Intransit Visibility (ITV), and the Total Distribution Advanced Technology Demonstration (TDATD) are individual, non-integrated solutions to the complex problem of physical distribution. The Defense Transportation System (DTS) establishes the process by which personnel and materiel move from origin to destination. The United States Transportation Command (USTC) has been established, and given unified command authority to develop a comprehensive transportation system for the DoD. The outline of USTC's approach to accomplishing this task is contained in the "Reengineering the Defense Transportation System Action Plan". A key provision of this concept is the empowered "DTS Agent", those persons who will ultimately be responsible for the efficiency and effectiveness of the DTS in peace, war, and Operations Other Than War (OOTW). Through the developing Global Transportation Network (GTN), USTC seeks a worldwide capability to meet customer requirements.

This paper proposes a methodology to address the Defense distribution problem based on an artificial intelligence, limited horizon approach, made possible through the integration of emerging technologies. By combining autonomous computational agent technology with the cost effective use of Low-Earth-Orbit (LEO) satellites for communications, a robust architecture capable of supporting worldwide power projection is achievable. The concept of decentralizing logistics control based on the use of autonomous computational agents was initially submitted to the Logistics Management Institute (LMI) by Dr. Michael N. Huhns, in response to a solicitation in the Commerce Business Daily. In the solicitation, LMI was seeking new and emerging technology applications to improve strategic mobility on behalf of the Director for Logistics, Joint Staff (J4). This proposal expands Dr. Huhns' concept and combines it with other technologies to address Defense distribution. One of the technological challenges in this proposal is the unique requirement for agents to learn from and adapt to other agents' decisions.

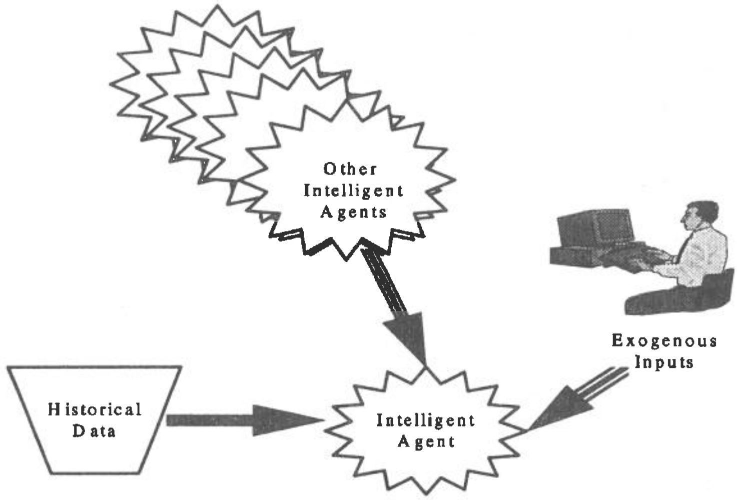

Figure 1

By combining these concepts our concept envisions cooperative Intelligent Agents (IA) strategically placed and integrated throughout the DTS [Gasser 1991]. Under this concept, the distribution network will be optimized sub-locally at each node using a limited look ahead based on stochastic, historical approximations of the global state of the system [Lee & Cohen 1985]. Each node in the distribution network and the shipment units, e.g. major weapons systems, containers, being shipped act as IA. The static IA (attached to facilities) must be capable of communicating with adjacent nodes and scheduling the transportation assets under their control. The mobile IA (attached to the items being shipped) must be capable of communicating with the static IA and negotiating their transportation according to their priorities and the available transportation assets. (See Figure 1 above.)

2. Description of IA

The IA act as expert systems acquiring new information and updating their decision making algorithms with each transaction. The IA make decisions based on the state of the system. The IA perception of the system comes from three sources: exogenous parameters, historical data, and information being passed in real time from other IA. During routine conditions, the overall goal of the static IA will be to deliver the materiel as needed using the most efficient transportation means, consistent with customer priorities. During crisis situations, the emphasis will shift from efficiency to maximizing the velocity and agility of the system. They become the drivers which enable an efficient, responsive and cost effective Defense Distribution System. Initially, the rules governing IA behavior will be extracted from the various regulations and directives that govern movement under the DTS. The system standards provide the foundation from which machine learning will occur. The agents function within the decision support system provided by USTC under their Command, Control, Communications and Computer Systems (C4S) architecture. This will eventually be the GTN.

From a macroeconomics standpoint, we can view each of the IA as a producer. Each producer attempts to maximize its utility, and given a sufficiently competitive system, the overall system achieves a stable, optimal macroscopic solution [Hahn 1982]. In the context of this problem sufficiently competitive means a relatively large number of shipments with varying priorities and adequate available transportation modes and assets at the distribution systems decision points. The shipment units' utility is inversely related to the time it takes them to reach their destination in a safe, undamaged condition. The static IAs' utility is measured in terms of minimizing cost under routine conditions and a function of shipment priorities and arrival times under crisis situations [Ferguson, et. al. 1989].

We could describe the objective of the mobile IA to be:

$$\underset{\forall Action}{Min} \{Arrival\ Date + M\phi\}$$

Equation 1

where ***Arrival Date*** is the date the shipping unit arrives at its destination, f represents the condition of the shipment (e.g. the smaller is the better), and ***M*** is the penalty for delivery of the shipment in less than perfect condition.

The static IA have an objective function which is dependent on both the cargo flowing through their areas of responsibility and the transportation assets the IA assign to that cargo. A sample objective function for a static IA is given below in Equation 2.

$$\underset{\forall Action}{Min}\left\{ \underbrace{\Pi \cdot \Psi}_{\text{Shipment Units}} + \underbrace{\beta\Phi(\bullet)}_{\text{Condition Violations}} + \gamma \cdot \underbrace{\sum_{\forall \text{Shipping Unit}} Cost}_{\text{Shipping Costs}} \right\}$$

where $\Pi \triangleq \{P_1, P_2, \cdots\}$ a vector of weightings (or priorities)

for the various shipments, $\Psi \triangleq \left\{ \begin{array}{c} \{Arrival\ Date + M\phi\}_1 \\ \{Arrival\ Date + M\phi\}_2 \\ \vdots \end{array} \right\}$,

$\beta \equiv$ weighting for the condition violations, $\Phi(\bullet) \equiv$ a function relating violations of shipping and storage conditions to a numerical value (e.g. storing ammunition at too great a density, sending refrigerated goods in non - refrigerated vehicles), $\gamma \equiv$ weighting for shipping costs, and Cost $\equiv$ cost of shipping from this node forward

Equation 2

Note in both the case of the static and the mobile IA, the IA will be making inferential decisions, i.e. the IA's objective functions are not fully determined until the shipments all reach their destinations, but the IA make decisions in real time while the shipment units are still in the transportation pipeline. Therefore, the IA are forced to make inferences about what the optimal decisions are from an objective function standpoint. The IA will base these decisions on the three sources of information already mentioned: exogenous parameters, historical data, and information being passed in real time from other IA. Of particular note here is the historical data; as IA complete transactions the historical data base will be updated, and future IA decisions will be modified based on these updates. For example, imagine a situation where a particular carrier is very unreliable. Every time an IA assigns a shipment unit to the carrier, the shipment unit is lost. Very quickly the system would start to avoid this carrier, because the expected value for the shipment unit cost would rise.

The mobile and static IA would vary in complexity as well as function. The mobile IA could be relatively simple. The mobile IA could carry a database consisting of few elements. For example, it might contain: a unique identifier, a pointer to the detailed contents (manifest), brief aggregate description of contents, current location, destination, priority and condition. The identifier is the distribution system's method of distinguishing that particular shipping unit. The shipping unit does not need to implicitly have all of its cargo identified in great detail, but it must be able to tell the user where the information is. The brief description helps human managers when intervention is required. Current location of the shipping unit and its destination are used to aid in routing decisions. The priority is the relative importance of the shipment. Condition covers special storage and cargo handling issues such as: explosive, requires refrigeration, fragile, etc.

The mobile IA are required to perform two tasks. First, they must be able to query and respond to the queries of other IA in the distribution system. Second, as rules surrounding their shipments are violated, the IA must identify those violations to decisions makers. For example, IA might complain if: the shipment needed to be refrigerated and it was too hot; it was fragile, and it was at the bottom of a stack; or the shipment has exceeded its required delivery date.

In addition to the tasks performed by the mobile IA, the static IA must be able to select modes for the various shipping units entering its zone of responsibility. Further, the static IA are responsible for maintaining the inferential data. Each time a transaction is completed, the static IA must refresh its associate database. The location of the various IA will be illustrated in the following paragraphs.

All requests for DTS shipments must initially be processed through a Transportation Office. This is the first organization in the transportation hierarchy that has the authority to commit Government funds for the movement of DoD materiel and personnel. The procedures are essentially the same for both peacetime and wartime shipments. We envision the information system in the Transportation Office as being the first IA. The purpose of these IA will be to facilitate the process of receiving requests for transportation and arranging for the movement of the associate cargo or passengers. Examples include: household goods, privately owned vehicles, general freight, individual and group moves. The IA would be expected to perform the following tasks: process routine requests for movement under established guidelines (priority of shipment, type of service, mode of transportation, least costs for requested service, etc.); develop the ability to gradually assist with more complicated shipment procedures (e.g. shipment consolidations, mode selection, etc.); interface with the next higher echelon of control for unit moves, with respect to granting approval of the shipment to begin moving (port call phasing to prevent port congestion); acquiring information about previous shipment movement performance for the purpose of optimizing future shipments (faster service, better carrier, less cost, etc.).

The next series of IA would be associated with the shipment units themselves. This capability is what was envisioned by [Huhns 1994], where the cargo would actually route itself through the distribution network. It will be useful to explore the problem by examining the transportation network in more detail. The Advanced Research Projects Agency (ARPA) has an initiative called TRANSTECH which will explore modeling the transportation infrastructure to include both information, transportation asset and cargo flow. One of the most important aspects of the DoD transportation system is that the components must be able to communicate world wide. The integration of intelligent agents with the use of low earth orbit (LEO) satellites for communications, appears to offer great military potential.

There are numerous options as to the level of shipment unit (piece, container, etc.) the agent(s) would be assigned. The IA associated with the cargo would be required to perform the following tasks: interface with the distribution system according to the priority established by higher authority; communicate status by LEO satellite to adjacent and higher echelons of control; function under established performance criteria, and notify appropriate control echelons when the criteria has not been met or has been exceeded (e.g. sent to the wrong port, in port too long, etc.); notify its consignee of its status. An example of how the IA could apply rules is shown in the following table.

LOCATION OF IA	STATE OF SYSTEM	TRANSITION ACTION	IA QUERY	IA TRIGGER	IA ACTION
Shipping Activity [Static IA]	Loading	Order Transportation	Is container moving?	Location of IA is within 1km of last location	Notify OPNS Center; carrier to pickup, and shipper
Intransit within CONUS {Mobile IA]	Movement		Is container at East or West Coast port within 4 days?	Location of IA compared to East or West coast port and number of days from stuffing	Notify OPNS Center; request carrier determine problem and notify shipper
Port [Static IA]	Cargo arrives at transportation node	Transfer to new mode	If location is a port and location static	Location of IA and number of days static	Notify OPNS Center and carrier
Intransit over ocean {Mobile IA]	Movement		Is container on schedule?	Location of container at least 432 NM from previous location at current time minus 24 hours	Notify OPNS Center, carrier, and "ship to" addressee
[Mobile and Static IA]		Divert cargo to new mode and destination	Is location an interchange point?	Location of IA is an interchange point (i.e. port)	Notify carrier to offload container and divert to new mode and destination; notify old and new "ship to" addressee; notify OPNS Center

Table 1

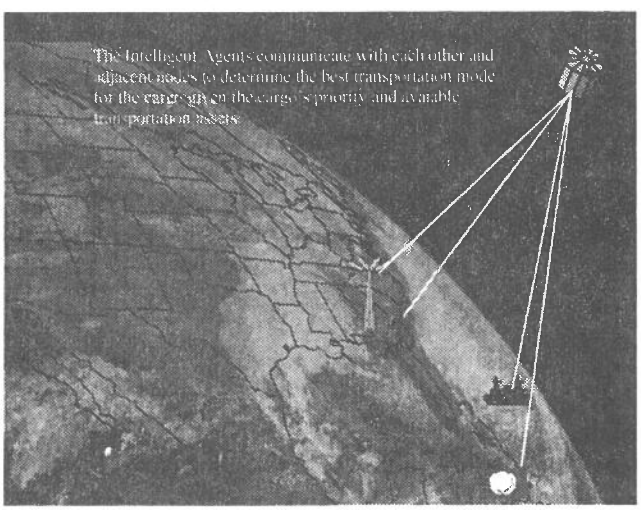

Figure 2

There are two primary types of intelligent agents, static and mobile. (See Figure 2 to the left.) The mobile intelligent agents are attached to shipments of cargo and passengers while the static intelligent agents are associated with the distribution nodes.

How the agents cooperate within the distribution network and compete for limited movement resources is at the center of our current research. Dr. Huhns suggested a method for achieving control and concomitantly efficiency, by making use of a market mechanism [Wellman 1992]. In this technique, each intelligent entity would be given a sum of "money" that it can use to purchase transportation, storage, or whatever other resources it might need. The amount of "money" available would vary with the priority of the shipment unit within the distribution network, and could change dynamically. An example might be specific repair parts moving through the systems as "general resupply" when they are elevated to "critical commodity" to support weapons system repair for pending operations. The key here, is that the LEO satellite network can function as an interrogator for the information about each shipment unit. The system: offers global coverage, can link players from around the world, and is therefore ideally suited to the military requirement to support force projection.

The next location where intelligent agent technology would be employed is at the distribution control level. This level of activity is difficult to define, because much of what is envisioned is currently only in the conceptual phases. As the designated CINC for transportation in both peace and war, USTRANSCOM is now developing the capability to perform its statutory functions in the 21st century. A vision of the future exists in USTRANSCOM's concept of "DTS 2010". This document describes the future environment, without describing specifically how it will be achieved. The use of IA can contribute significantly to developing the required capabilities for the next century. The first task facing USTRANSCOM is to integrate data from numerous sources into an effective decision support system.

This task will begin with the development of GTN. It should become the command and control system for the DTS. The challenge facing GTN development is the meaningful integration of data from a variety of sources, both within and external to the Department of Defense (DoD). USTRANSCOM is composed of a headquarters and three component commands; the Military Traffic Management Command (MTMC) (land transportation and sea port operations); the Military Sealift Command (MSC) (sealift operations); and the Air Mobility Command (AMC) (airlift operations). Each of these organizations interface with each other and each relies on the commercial sector to provide transportation services. The problem remaining is the vertical integration of data for each of the component commands and their many supporting automated systems and their association with the private sector and their suite of automated systems. Likewise, the horizontal integration of data among the component commands remains a challenge. Some components have complex systems, such as AMC, while some are relatively unsophisticated, as MSC, while MTMC deals almost exclusively with the commercial sector. USTRANSCOM headquarters, and the GTN prototype, have made some effort toward collecting data from a variety of sources and displaying it. The thrust as been to gain intransit visibility, without developing the underlying management structure. Currently, there is no effort underway to begin linking all the data into a comprehensive decision support system. Here, we believe, is the potential for the greatest payoff.

IA would be developed to remotely monitor the real-time performance of the DTS at the shipment unit level and alert the decision makers when problems occur so corrective action can be taken. Data is currently available from a variety of sources; DoD owned systems as well as commercial systems which support the movement of military equipment and supplies. IA would be expected to: monitor cargo from origin to destination (using LEO satellites), and alert the human managers when expected performance criteria are not met; avoid port congestion, and automatically hold or re-route cargo; recommend cargo routing based on the availability of lift assets and there schedules; recommend how to accomplish cargo diversion at the least cost using the fastest method; anticipate and react to problems; and interface with the IA at the transportation office and with the cargo to facilitate the efficiency of movement.

At this level, the IA act as the network controllers. They are interested in the overall performance of the system from origin to destination. As described, these IA are helped along the way by other agents.

The final series of supporting agents concern themselves with ensuring that the cargo reaches its final destination (referred to in the military as the "theater"). The locations of these agents is not static, but rather they are placed to support the concept of force projection. They must accompany a force physically or provide support electronically. These IA would be expected to keep track of unit locations and assist incoming cargo to find its way to intended recipients. The IA would also: recommend cargo clearance actions based on priorities and the distribution scheme in theater; respond to other IA concerning the ability to move cargo in theater; assist with optimizing available in theater lift assets to meet expected cargo arrival volumes; recommend the most effective manner to accomplish a cargo diversion; generally manage distribution. Inherent in this requirement is the linkage between the supply, maintenance, and transportation systems. Collectively, they comprise the distribution function. As our concept matures, its principles can be adapted to assist in managing both theater supplies and transportation assets. The theater system is nothing more than the strategic system, only confined to a smaller physical area. The process and procedures are exactly the same.

We believe it possible to employ intelligent agent technology to integrate the distribution process from origin to destination in a manner which ensures synergy.

References

[Ferguson et al. 1989] D. Ferguson, Y. Yemini, and C. Nikalson, "Microeconomic Algorithms for Load Balancing in Distributed Computer Systems," Research Report, IBM T.J. Watson Research Center, Yorktown Heights, NY, October 1989.

[Gasser 1991] Les Gasser, "Social Conceptions of Knowledge and Action: DAI Foundations and Open Systems," Artificial Intelligence, Vol. 47, pp. 107-138, 1991.

[Hahn 1982] F. Hahn, "Stability," in K. Arrow and M. Intriligator, eds., Handbook of Mathematical Economics II, Chapter 16, North Holland Publishing Co., 1982.

[Hewitt 1991] Carl Hewitt, "Open Information Systems Semantics for Distributed Artificial Intelligence," Artificial Intelligence, Vol. 47, pp. 79-106, 1991.

[Hewitt & Inman 1991] Carl Hewitt and Jeff Inman, "DAI Betwixt and Between: From `Intelligent Agents' to Open Systems Science," IEEE Transactions on Systems, Man, and Cybernetics, Vol. 21, No. 6, pp. 1409-1419, November/December 1991.

[Huhns 1994] Michael N. Huhns, "Decentralized Logistics Using Autonomous Agent Technology," (unpublished white paper), Microelectronics and Computer Technology Corporation, Austin, TX, November 1994.

[Huhns & Bridgeland 1991] Michael N. Huhns and David M. Bridgeland, "Multiagent Truth Maintenance," IEEE Transactions on Systems, Man, and Cybernetics, Vol. 21, No. 6, pp. 1437--1445, November/December 1991.

[Kephart et al. 1989] J. O. Kephart, T. Hogg, and B. A. Huberman, "Dynamics of Computational Ecosystems: Implications for DAI," in Les Gasser and Michael N. Huhns, eds., Distributed Artificial Intelligence, vol. II, Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1989.

[Lee & Cohen 1985] L.L. Lee and M.A. Cohen, "Multi-Agent Customer Allocation in a Stochastic Service System," Management Science, vol. 31, pp. 752-763, June 1985.

[Singh & Huhns 1994] Munindar P. Singh and Michael N. Huhns, "Automating Workflows for Service Order Processing: Integrating AI and Database Technologies," IEEE Expert, October 1994.

[Wellman 1992] Michael P. Wellman, "A General Equilibrium Approach to Distributed Transportation Planning," Proceedings AAAI-92, San Jose, CA, July 1992, pp. 282-289.

Mutually Supervised Learning in Multiagent Systems

Claudia V. Goldman and Jeffrey S. Rosenschein
Computer Science Department
Hebrew University
Givat Ram, Jerusalem, Israel
ph: 011-972-2-658-5353
fax: 011-972-2-658-5439
email: clag@cs.huji.ac.il, jeff@cs.huji.ac.il

Abstract

Learning in a multiagent environment can help agents improve their performance. Agents, in meeting with others, can learn about the partner's knowledge and strategic behavior. Agents that operate in dynamic environments could react to unexpected events by generalizing what they have learned during a training stage.

In this paper, we propose several learning rules for agents in a multiagent environment. Each agent acts as the teacher of its partner. The agents are trained by receiving examples from a sample space; they then go through a generalization step during which they have to apply the concept they have learned from their instructor.

Agents that learn from each other can sometimes avoid repeatedly coordinating their actions from scratch for similar problems. They will sometimes be able to avoid communication at run-time, by using learned coordination concepts.

Subtopic: Learning in multiagent systems, Coordination

1 Introduction

Distributed Artificial Intelligence (DAI) is concerned with effective agent interactions, and the mechanisms by which these interactions can be achieved. One of the central issues in multiagent environments is that of appropriate coordination techniques. Much DAI research deals with this issue by giving pre-computed solutions to specific problems. For every new problem, the agents will start from scratch and derive the appropriate solution (even if it is an interaction instance identical to one they have seen before). For example, researchers have considered negotiation as a technique for deriving agreements that determine agent actions [Smith, 1978; Durfee, 1988; Kraus and Wilkenfeld, 1991; Zlotkin and Rosenschein, 1993]. These negotiation techniques invariably focus on a single encounter or set of encounters; agents do not (for example) improve their negotiation performance based on experience. Other DAI researchers have focused more on direct modeling of agents' beliefs and desires, as another way for an agent to decide what action to perform when dealing with others [Grosz and Kraus, 1993]. Again, learning rarely enters into this research; while the exploitation of a model of the opponent is studied, the actual derivation of the model rarely is.

Multiagent reactive systems have also been analyzed within DAI, where solutions are arrived at dynamically by reactive agents in multiagent environments. Social laws [Tennenholtz and Moses, 1989; Shoham and Tennenholtz, 1992] and cooperative state-changing rules [Goldman and Rosenschein, 1994] have been studied; these conventions give the agents a framework within which to act, to more harmoniously interact with the other agents participating in the same world. Learning has been investigated within this framework, particularly in [Shoham and Tennenholtz, 1994], which investigated how conventions can evolve when the Highest Cumulative Reward update rule is used (i.e., agents choose to perform the action that has yielded the highest payoff until then).

The advantages of having agents learn within a multiagent environment are clear. In Cooperative Problem Solving systems, cooperative behavior can be made more efficient when agents adapt to information about the environment and about their partners. In competitive Multiagent Systems, agents' performance within the environment can be improved if they can learn about the strategies and preferences of their opponents.

In this paper, we present a learning algorithm for a cooperative multiagent environment. The agents in our model first go through a training step, and are then able to choose their actions by generalizing what they have learned. The agents do not need to re-coordinate their actions for every new situation or problem. The main issue in our research is how to train the agents in a way that minimizes the number of mistakes in the generalization step.

This distributed learning approach to coordination is useful whenever agents do not have enough time to negotiate, when they exist in a dynamic environment and will benefit by adapting to unpredictable situations, and when the agents face similar problems repeatedly. The actual utility of having agents use learning instead of, for example, negotiation, will depend on the tradeoff between the time it takes to train an agent (which induces

Organizing Committee

Sandip Sen, University of Tulsa, Chair
Michael N. Huhns, MCC
Maja Mataric, Brandeis University
Michael Shaw, University of Illinois at Urbana-Champaign
Devika Subramanian, Cornell University
Ming Tan, GTE Laboratories Inc.
Moshe Tennenholtz, Technion - Israel Institute of Technology, Israel
Gerhard Weiss, Technische Universitaet Muenchen, Germany

MICHAEL N. HUHNS

IJCAI–95 Workshop Program
Working Notes

Adaptation and Learning in Multiagent Systems

Monday, August 21, 1995

Sponsored by the International Joint Conferences on Artificial Intelligence Inc. (IJCAII), American Association for Artificial Intelligence (AAAI), and the Canadian Society for Computational Studies of Intelligence (CSCSI)